

Running head: Children and robots learning to play hide and seek

Children and robots learning to play hide and seek

J. Gregory Trafton, Alan C. Schultz, Dennis Perzanowski, Magdalena D. Bugajska

William Adams, Nicholas L. Cassimatis, and Derek P. Brock

Naval Research Laboratory

Abstract

How do children learn how to play hide and seek? At age 3-4, children do not typically have perspective taking ability, so their hiding ability should be extremely limited. We show through a case study that a 3 1/2 year old child can, in fact, play a credible game of hide and seek, even though she does not seem to have perspective taking ability. We propose that children are able to learn how to play hide and seek by learning the features and relations of objects (e.g., containment, under) and use that information to play a credible game of hide and seek. We model this hypothesis within the ACT-R cognitive architecture and put the model on a robot, which is able to mimic the child's hiding behavior. We also take the “hiding” model and use it as the basis for a “seeking” model. We suggest that using the same representations and procedures that a person uses allows better interaction between the human and robotic system.

Children and robots learning to play hide and seek

Introduction

There are, of course, many ways to build a computational system that behaves intelligently and works well with people. Our working hypothesis, the representational hypothesis, is that a system that uses representations and processes or algorithms similar to a person's will be able to collaborate with a person better than a computational system that does not. While we believe that our hypothesis is quite general, we will focus the majority of the paper on robotic systems. There are, of course, many ways of interacting with robots (using a joystick or similar device is one of the most common), but in order to have full collaboration with an intelligent system, the person and computational system need to communicate with each other. We focus on physical robots because we have a strong belief that a system that has sensors and effectors (e.g., embodied cognition) is a first step to achieving strong collaboration with a person. We suggest three reasons for the representational hypothesis and then describe empirical and computational evidence in the domain of the children's game hide and seek.

First, since algorithms written for traditional real-time robotic systems have to be computationally efficient, they tend to use efficient mathematical representations such as matrices and polar coordinates, which may not be natural for people to use. For example, most position and motion information in robotics is conveyed using position vectors and transformation and rotation matrices. In general, people do not think or reason in this format. Instead, people seem to use a combination of spatial and propositional knowledge (Anderson, Conrad, & Corbett, 1989; Anderson & Lebiere, 1998; Shepard & Metzler, 1971; Taylor, 1992; Trafton et al., 2000; Trickett, Ratwani, & Trafton, under review). Thus, in order to interact with

the robot, the system must translate the robot's representation to a person's representation. Because the person's representation of space is so complex (Harrison & Schunn, 2002, 2003a, 2003b; Previc, 1998), this is not a trivial task. Additionally, a translator does not allow shared operations to occur between the person and the system; all operations must go through the translator, which may cause some loss of information, or confusion to either or both systems.

Second, if a human is going to collaborate in shared space with a robot, the robot should not exhibit unexpected, unnatural or “martian” behaviors (Petty, 2001). While the robot may be able to efficiently perform a task using, for example, a behavior-based approach, if the resulting behavior is perceived to be unnatural by the human, it will detract from the interaction. Therefore, we propose that some robot behaviors be created by modeling how humans perform such a task, and then using that model to drive the robots behavior.

Third, and most important to this paper, is that we believe that some tasks for robots can best be programmed not by using more traditional control algorithms, but through the understanding of how humans solve the task. So if, for example, we want to create a robot that can search for hidden snipers, it makes sense to encode knowledge about how humans hide. We believe this can be best achieved with computational cognitive models.

In this project we seek to understand how children learn to play hide and seek, and thus create a robot that understands how to play hide and seek. We have chosen hide and seek because it ties in well to several of our goals. Hide and seek forces us to work in a complex, dynamic environment, it allows us to explore embodied cognition issues (i.e., spatial and temporal reasoning and allows us to explore methods of combining both cognitive and robotic/AI methods into a single system. Hide and seek also encapsulates many interesting real-world tasks

including urban search and rescue, homeland defense (finding nuclear, biological, or chemical weapons), and military actions such as finding hidden troops or locating the best areas in which to be concealed.

For the remainder of the paper, we will describe the target robot, its sensors, navigational system, and how it communicates with people. Next, we describe the cognitive question we are investigating, describe a case study, and describe the computational cognitive model and how it operates on the robot. Finally, we will directly explore our hypothesis via computational means by examining how well the system can generalize to other tasks.

Robot System

This section describes the robot hardware and software.

Hardware

The robot is a commercial Nomadic Technologies Nomad200 suited to operation in interior environments. It has a zero turn radius drive system, an array of range, image, and tactile sensors, and an onboard network of Linux and Windows computers with a wireless Ethernet link to the external computer network.

Software

A combination of non-cognitive methods (primarily for robot mobility and object recognition), cognitively-inspired interactions (primarily for communicating with a person), and computational cognitive models (primarily for the high-level thinking and reasoning) were used. We have previously shown the utility of combining low-level reactive systems with cognitive models (Bugajska, Schultz, Trafton, Mintz, & Gittens, 2001; Bugajska, Schultz, Trafton, Taylor, & Mintz, 2002; Trafton, Schultz, Bugajska, Gittens, & Mintz, 2001).

Non-cognitive Methods:

This project draws on the robot mobility capabilities of the previously developed WAX system (Schultz, Adams, & Yamauchi, 1999), which includes components for map building, self localization, path planning, collision avoidance, and on-line map adaptation in changing environments. The robot's lowest level of information comes from a dead-reckoning component that integrates motion over time to compute the robot's current location. As the robot moves, it gathers range data from its 16 ultrasonic transducers and a laser-based structured light rangefinder. In a process developed by (Moravec & Elfes, 1985), the range data is interpreted using a sensor model that converts the raw range to a set of occupancy probabilities for the sensed area. In this manner, data from multiple sensors can be fused into a single short-term occupancy map of the robot's vicinity, represented as a three dimensional array of discrete cells, each containing the probability that it is occupied or empty.

All robots' odometry suffers from gradual drift, sometimes punctuated by larger errors from wheel slippage, rough ground, or collisions, so odometry alone is insufficient. Using the process of continuous localization (CL) (Schultz & Adams, 1998), a temporally overlapping progression of short-term maps is maintained. At periodic intervals, the oldest short-term map, which has the most data, is registered against a long-term map of the larger environment (typically a room) to determine the correction needed to correct the odometric drift. The long term map can be supplied a priori, or learned through a careful exploration as was done in Yamauchi, Schultz, & Adams (1998). For this work, mapping was not the focus so an a priori map was used. As a byproduct of correcting odometry, the long-term map can also be adapted to incorporate the now-corrected new readings from the short-term map. Thus, as the robot

moves, it not only maintains an accurate estimate of its position but also keeps the long-term map up to date with any changes to the environment.

Because the robot's basic motor system is geometry-based and metric maps can be easily produced, it is a matter of practicality to state goal locations as points in Cartesian space. These goals are passed to the Trulla path planner (Hughes, Tokuta, & Ranganathan, 1992), which uses the long-term map to determine the best path to the goal. For a given goal and map, planning begins at the goal and works outward. Each neighboring map cell is assigned a vector pointing to a neighbor that has the least cost path to the goal so far. This process is recursive, and all cells are visited. When exhausted, each map cell contains a vector pointing in the direction of the least cost path to the goal, free of any local minima (though sometimes inadequacies in the conversion from occupancy probability to traversability by a non-point robot can result in non-traversable paths).

Because there may have been changes to the environment that are beyond the robot's sensor range, or recent changes such as people walking near the robot, the paths made by Trulla cannot be followed blindly. Instead, they are passed as a single vector field to the Vector Field Histogram (VFH) process (Borenstein & Koren, 1991). VFH uses the robot's current position to retrieve from the vector field the direction the robot should move to head toward the goal. This vector is compared to an occupancy histogram built from the short-term map (which has the recent data close to the robot) and the robot is steered in the unblocked direction closest to that indicated by the vector. In effect, Trulla handles the room-level navigation while VFH provides collision avoidance. If the robot is blocked, VFH prevents collision, CL learns the changes and produces a new adapted long-term map, and Trulla replans around the obstruction.

Rather than providing the robot with a priori information about discrete objects for it to hide behind, the robot was instead equipped with limited computer vision in order to detect some objects autonomously. This also allows objects to be rearranged, added, or removed with the robot reacting accordingly. The CMVision package (Bruce, Balch, & Veloso, 2000) was used to provide simple color blob detection using an inexpensive digital camera mounted on the robot.

Objects are tagged with a special color marker that is more easily distinguished from the surroundings. The marker color is the identifier for the characteristics of an object. For example, all lime green objects are "chairs" and have the same characteristics. A table is supplied that maps marker color to the object's size, but all information on hidability is learned through feedback from playing the game and added to the table to be used in subsequent games. The bearing to the object is then determined from its location in the camera image, and the range to it is obtained from a scanning laser rangefinder.

Cognitively Inspired Methods:

In order to communicate with a person, we use several methods that have some basis of human cognition. The methods that are used here allow a user to communicate with the robot using spoken language, gestures to the robot, and gestures on a Palm PDA.

The human user can interact with the mobile robot, using natural language and gestures, which are part of our multimodal interface. The natural language component of the interface uses a commercial off-the-shelf speech recognition engine, ViaVoice, to analyze spoken utterances. The speech signal is translated to a text string that is further analyzed by our in-house natural language understanding system, Nautilus, to produce a regularized expression. This latter

representation is linked, where necessary, to gesture information, and an appropriate robot action or response results.

For example, the human user can tell the robot “Coyote, go hide and I’ll try to find you.”

The speech signal is analyzed into a text string which when parsed produces the following representation, simplified here for expository purposes.

```
(and (imperative (p-hide: hide)
                  (system: you
                    (name: coyote)))
      (future (p-attempt: try)
              (agent: I)
              (action (p-find: find)
                      (agent: I)
                      (system: you
                        (name: coyote))))))
```

Basically, Nautilus parses the utterance into appropriate commands (e.g. the *imperative* structure in our example) and statements (e.g. the *future* declaration in our example), and the various verbs or predicates of the utterance (e.g. *hide*, *try*, and *find*) are mapped into corresponding semantic classes (*p-hide*, *p-attempt*, and *p-find*) that have particular argument structures (*agent*, *system*) which result in a semantic interpretation of the utterance. With gesture information, where appropriate, these representations are then sent to the robotic component whose modules translate these representations into appropriate actions.

In the example above, no further gesture information is required to complete the command. Coyote will, therefore, respond “I will go and hide,” in order to inform the user that it has understood the utterance, and the appropriate behavior based on the cognitive model for the hide-and-seek activity is invoked and appropriate robot action according to the model ensues. If, for example, a gesture is required to disambiguate the speech, as in “Coyote, hide somewhere over there,” then gesture information obtained from the laser rangefinder mounted on the top of the robot indicates the desired location, and this information is included in the interpreted utterance for further analysis by the robotic system.

A more detailed analysis of how our multimodal interface processes both natural language and gestures, mapping them to appropriate robot actions and responses, is available elsewhere (Perzanowski, Schultz, & Adams, 1998; Perzanowski et al., 2002; Perzanowski, Schultz, Adams, & Marsh, 2000; Perzanowski, Schultz, Adams, Marsh, & Bugajska, 2001; Skubic, Perzanowski, Blisard, Schultz, & Adams, in press).

Hide and Seek

We are exploring our representational hypothesis within the children's game hide and seek. Hide and seek is a simple children's game where one child is "It," stays in one place to count to ten, and then goes to seek, or find, the other child or children. These issues address our high-level goals of understanding how humans represent and process spatial information, particularly as an aid in designing better human-robot interaction in collaborative spaces. Our specific goal in this study is to understand how children learn to play hide and seek and to use this knowledge to build a computational cognitive model to enable a robot to play hide and seek

with near human level decision making (or competence). Our cognitive model was written in ACT-R (Anderson & Lebiere, 1998).

How do children learn how to play hide and seek? Specifically, how do children learn how to hide? Young children can play peek-a-boo at approximately 7 months of age (Kleeman, 1973) as they are just developing object permanence, shown to begin somewhere between five months (Baillargeon & DeVos, 1991; Baillargeon, Spelke, & Wasserman, 1985; Bower, Broughton, & Moore, 1971) and nine months (Piaget, 1954)

However, a "good" hider needs spatial perspective taking to be able to find the best hiding places. For example, a good hider must take into account where "It" will come into a room, where "It" will search first, and where to hide behind an object from the perspective of "It" (e.g., Lee & Gamard, 2003). A good hider also needs to know that just because the hider can't see "It" doesn't mean that "It" can't see the hider. Finally, keeping an object (like a column) in between "It" and the hider is frequently a good hiding tactic. All of these issues need some form of spatial perspective taking ability, or the ability to see the world from someone else's eyes.

Children begin to develop very rudimentary spatial perspective-taking ability around age four (Huttenlocher & Kubicek, 1979; Newcombe & Huttenlocher, 1992; Wallace, Allan, & Tribol, 2001). In fact, Piaget and Inhelder (1948) claimed that even children eight or nine years old did not have perspective taking ability.

Previous researchers have studied perspective taking ability by examining children's egocentrism (e.g., Flavell, Omonson, & Latham, 1978) or spatial perspective taking (e.g., Newcombe & Huttenlocher, 1992). For example, a common methodology (based on Newcombe & Huttenlocher, 1992) is to bring a child into the laboratory and show them a table that has four

chairs with different objects or scenes visible from each chair (a desk, a window, etc.). The child sits down at one of the chairs while the experimenter sits down at another chair. The child is then asked to either describe or pick out from a set of pictures what the child sees (no perspective taking needed) and what the experimenter sees (spatial perspective taking needed).

This line of research has shown that four year olds have rudimentary spatial perspective taking ability in this kind of situation: 67% of the time, four year olds made correct "near-far" perspective taking decisions (Newcombe & Huttenlocher, 1992, experiment 2). However, four year olds did not seem to be able to differentiate "left-right" perspectives (Newcombe & Huttenlocher, 1992, experiment 2). These experiments have been replicated and extended (e.g., Wallace et al., 2001); the basic finding seems to be that four year olds have some very rudimentary spatial perspective taking ability, but it is nowhere close to a full understanding about how other people see the world differently (i.e., even the near-far accuracy, 67%, while better than chance could not be said to be "good" performance).

Additionally, hide and seek seems to be rather more complicated than some of the simple tasks that have been explored in the laboratory. Hide and seek typically occurs in a large-scale environment where the child can not see the entire area at once. Also, "It" may come into a room where the child is hiding in different ways (i.e., from different doorways), and the hider needs to determine if an object is big enough to get inside of or hide behind. Finally, there are many other things that are part of the game, including time pressure, the large number of available objects, and the large number of rooms or alternate locations.

Thus, according to this analysis, children under four should not be able to play a credible game of hide and seek. There doesn't seem to be any empirical investigations of the naturalistic

game of hide and seek, but large amounts of anecdotal evidence (i.e., a casual examination of the game-playing behavior at local parks) suggests that even three year old children can certainly play a credible game of hide and seek. If even four-year old children do not have a good model of spatial perspective taking, how do they learn to play hide and seek?

Our hypothesis is that, since perspective taking is not learned well until later, a child before four or four and a half will not be able to use spatial perspective taking as a primary strategy in the hide and seek game. In order to play the game, three and four year old children instead learn features and relations of objects (e.g., Smith, 2000) that are pertinent for the hiding game. For example, they need to learn that whether or not it is possible to see through an object is an important feature (opaque/transparency). Likewise, they need to learn that size is an important feature (i.e., is an object big enough to get inside of). Thus, the relationship of different aspects of objects is the key. One implication of this hypothesis is that hiding behind something will be a rare occurrence because hiding behind something requires some perspective taking ability. If hiding behind an object does occur, it will probably occur only in a very familiar environment.

In order to investigate this object-relationship hypothesis, we collected data from a single child, aged 3 1/2, who was learning to play hide and seek. We then built a computational cognitive model of the hiding behavior seen by the child in the case study. We put this model on our mobile robot to see if we could get reasonable human-level decision making. Finally, we show support for our representational hypothesis by reusing our hiding model as the basis for a seeking model.

Case Study

Participant

The child used in the case study is the daughter of one of the authors of the paper. At 3 1/2, the child, E, did not know how to play hide and seek and needed the rules to be explained to her.

Task and materials

Fifteen games of hide and seek were played over a 4 hour period, with one break. Four games were played first, then a break occurred, and then the final eleven games were played later in the day. "It" counted to ten while E hid. The game occurred inside E's house, and E could hide anywhere in the house. E was, of course, very familiar with the house. "It" had a video camera on the entire time a hide and seek game was played, recording the interactions between "It" and E as well as the final hiding places that E chose. In one game (not included in the fifteen above), the roles were switched: E was "It" and "It" hid.

Over the following few days, the spatial perspective-taking of E was examined by asking her to name her own left and right hands and other people's left and right hands. When asked to name other people's left and right hands, the other person was either sitting in the same direction as E or facing E.

Design and Procedure

"It" searched for E after "It" had counted to ten. In one case (described below), "It" provided E with a vague hint. In all other cases, "It" gave E some sort of feedback, ranging from "That's a better hiding place!" (positive feedback) to "I can still see you!" (slightly negative feedback).

Measures

All verbal utterances were transcribed and all but the first two hiding places were coded according to the type of hiding place it was. The first two games are described later. The specific codes we used were Under (E hid directly under an object), Containment (E hid inside of another object), and behind (E hid behind an object from any perspective).

Caveats to the case study

The results of this case study should be taken with care. The nature of case studies is that there is only one participant or observation, and it is difficult to determine how generalizable the results are based on one study. This study may be even more susceptible to that concern because the child chosen for the case study was one of the authors' children (following Piaget) and in a familiar environment.

However, it is possible to perform an in-depth analysis of one participant that is sometimes more difficult or impossible to do in more traditional experimental settings. The focus in this study was on modeling the individual behavior at a fine grain level to lead to generalizations that could be tested computationally and empirically. A similar methodology has been used by many researchers in cognitive science (e.g., Anzai, 1979; Lovett, Daily, & Reder, 2000; VanLehn, 1991).

Results and Discussion

E clearly did not have full perspective taking ability: she could correctly name her own left and right hand, and anyone else's left and right hands if they were sitting in the same orientation. However, if E was asked to name a person's left or right hand while facing that person, she was less than 50% accurate, showing an egocentric bias.

If our object-relationship hypothesis is correct, we would expect to observe very few (if any) instances of E hiding behind objects. Instead, we would expect to observe a predominance of hiding under objects and inside of objects or rooms (containment). As **Table 1** shows, we found strong support for our hypothesis.

Game Number	Hiding Location	Hiding Type
1	eyes-closed	can't see me if I can't see you
2	out-in-open	understanding rules of game
suggestion	don't hide out in the open	
3	under piano	Under
4	in laundry room	Containment (room)
break		
5	under piano	Under
6	in laundry room	Containment (room)
7	in bathroom	Containment (room)
8	in her room	Containment (room)
9	under chair	Under
10	behind bedroom door	Containment or behind
11	under chair	Under
12	under covers	Under or containment
13	under covers	Under or containment
14	in bathroom	Containment
15	under glass coffee table	Under

Table 1: Summary of where E hid

The majority of places that E hid in were either containment (i.e., hiding inside a room) or under an object, or both (80%), though there was one instance where E hid behind a door (7%), one instance of hiding her eyes while out in the open (7%), and one instance of hiding out in the open (7%). Clearly, E was not using perspective taking skills to hide behind objects frequently, $\chi^2(2)=24.2$, $p < .001$, bonferonni adjusted χ^2 , $p < .01$.

For game #1, E went into a different room and closed her eyes, presumably thinking "If I can't see you, you can't see me." For game #2, E peeked at "It" around a corner as the counting was completed.

After game #2, "It" thought that E was stuck in a local minima, so he gave her a suggestion, "You might not want to hide in the open." This suggestion gave E a chance to think about the game a moment and come up with a new representation of the game of hide and seek. Immediately after this suggestion, E was able to dramatically improve her hiding behavior: she hid underneath a grand piano. She was still immediately visible when "It" came into the room, but she was doing more than simply hiding her eyes, and she was clearly not hiding out in the open.

For the next several games, E hid under things and inside of rooms. At game #9, she hid in what was probably the best hiding place of the entire day: underneath an upholstered chair. In this case, she was completely hidden from view from all angles in the room.

For the last few games, E explored other places, focusing primarily on hiding under things or inside of things. Note that some of the hiding places E used were ambiguous: hiding under bedcovers could be either a containment location (surrounded on all sides by the covers or an

under location (underneath the covers). Additionally, the only "behind" location was squeezed in between a closed door and the wall. This could be either a behind location or a containment location.

Several comments should be made about E's hiding places. First, it should be noted that E did not hide in the same room as "It" a single time. Second, for the first game (E hiding her eyes so she could not see "It", presumably thinking "If I can't see you, you can't see me") is strong evidence for E not having a well developed sense of spatial perspective taking. At this stage E did not completely understand the rules of the game, but she did understand that "It" should not be able to find her easily. If E had a well developed sense of spatial perspective taking, she would not ever have simply covered her eyes. Further, E chose the same hiding place multiple times. For example, five locations were used twice (under the piano, in the laundry room, in the bathroom, under the covers, under the chair). Also, after game #5, her hiding behavior gets markedly better --- in nine of the ten games after #5, she can't be easily seen. Finally, it appears that E understands at about game #4 that it is good to hide under things or within things (like small rooms). However, as game #15 shows, she does not yet understand that opacity is also a critical feature in this domain.

These kinds of hiding places strongly suggest that E is developing knowledge about objects and relations to objects in order to hide: she is probably not using spatial perspective taking in order to hide. E did not have spatial perspective taking ability measured by her ability to tell someone else's left or right hands in an orientation different from her own. She did not hide in places that would have shown spatial perspective taking (e.g., behind objects). However, her hiding places were quite good, especially after she had played several games. Can this type

of hiding behavior be modeled without spatial perspective taking? The next section examined this issue directly.

ACT-R Model

This is a very challenging task to model for several reasons. First, the learning that occurs happens very quickly and in very few trials. Second, there is a time limit to what kinds of hiding places can be found --- approximately 10 seconds to find a place and make the physical movements to the hiding place. Third, the model must be able to accept a suggestion and reason about that suggestion to change its behavior (i.e., get out of a local minima). Fourth, the model must be able to take positive or negative feedback and use that feedback to change its behavior. There is, in short, an enormous amount of learning that occurs in these 15 games with only one suggestion to the system.

We modeled this task in ACT-R (Anderson & Lebiere, 1998). The ACT family of theories has a long history of integrating and organizing psychological data (e.g., Altmann & Trafton, 2002; Anderson, 1983, 2002; Anderson, Bothell, Lebiere, & Matessa, 1998). The current version, ACT-R, derives important constraints from asking what cognitive processes are adaptive given the statistical structure of the environment (Anderson, 1990; Anderson & Milson, 1989). It has also been broadly tested in psychological and computational terms.

In order to learn and improve within hide and seek, several types of learning were used, including learning new knowledge structures (chunks) and schematic / ontological knowledge (links between these chunks), tuning of production rules, and a scaled down form of explanation based learning (Carbonell, Knoblock, & Minton, 1991). Our model focuses primarily on the first

few games, up to the point where E successfully uses knowledge of *containment* and *under* to find good hiding places. Our model successfully reasons with a suggestion provided to it.

The model begins every game by "examining" the environment. In the pure model (i.e., without the robot sensors), the model has environmental chunks that it can "see."¹ It starts off with a few specific hiding productions (based primarily on "peek a boo"), a few general reasoning productions, and a fair number of chunks and knowledge about the physical world. It also has some declarative knowledge about space -- knowledge about what "under" and "inside" means.

The model begins with very little a priori knowledge about how to hide. When asked to play hide and seek the first few times, the only strategy it has that is applicable is to close its eyes. Like E, the model is stuck in a local minima. In order to get out of the local minima, it needs some sort of suggestion or additional information. Again, like E, the model is told, "Don't hide out in the open." The model then reasons about what "open" means by examining the environment and reasoning explicitly about those objects. Specifically, it focuses its attention on an object (like a chair) and marks certain object-locations as "not-out-in-the-open." At the beginning of a model run, it believes that "under," "inside," and "on-top-of" are not-out-in-the-open. The model is then able to use that information (in competition with other hiding productions, like "hide-eyes") to find better hiding places the next time it is asked to play hide and seek.

¹ It should be noted that ACT-R/PM (Byrne & Anderson, 1998) has a mechanism for seeing the world. However, since the model needed to be able to transition to a robot with different sensor types (e.g., sonar), the model used environmental chunks rather than visual PM chunks. This simplification allowed the pure model form to solve the hide and seek problem rather than the vision problem.

Thus, the next time the model is asked to play hide and seek, it examines the environment and chooses at random a location that is "not out in the open", finds an appropriate object, and goes there. For example, the model is able to hide "under" a "piano," just like E did in game #3. In the current version, we provide the model with feedback (positive or negative) on every trial. In this case, the feedback would be negative, and the model would try a different location. Over several games (1-4), it is able to determine that some locations are better than others --- hiding inside of something is better than hiding on top of something. Within 4 games, the model is able to hide in reasonably good hiding places. At this point in time however, it does not know anything about transparency or opacity --- it is perfectly happy to hide under a clear glass coffee table, just as E did in game #15.

There are several interesting situations that arise in the model. As was noted earlier, E hid in the same place several times. The model shows the same pattern. The reason this seems to happen in the model is that when the model is "searching" for applicable objects, it is more likely to retrieve an object that has already been used because it has a higher base-level activation: it is more active or "hotter" in memory. We do not believe that the perceptual system works the same way that memory does, but after objects have been perceived, these objects may be subject to changes in activation even if they are in plain view. Thus, some objects and locations could become "favorite" hiding places simply because of the increase in activation. Because there is noise in the cognitive system, sometimes an object/location will be chosen multiple times and sometimes a different object/location will be chosen. This noise is one of the ways that ACT-R does not get stuck perseverating on the same objects and locations (Altmann & Trafton, 2002).

Additionally, the model is able to imitate E's hiding behavior quite well. Because there is randomness in the model, the initial performance of the model does not fit perfectly: the model may learn faster or slower than E did. However, with some guiding or model tracing of the model, it is able to perfectly match E's qualitative hiding behavior.

Finally, each time the model is asked to hide, it is able to find a hiding location within two to six seconds of ACT-R simulated time. Thus, there is approximately four to eight seconds to actually move to the hiding place. The model is therefore able to find a hiding place within the 10-second time limit set by the game.

Robot Behavior: Hiding

Our next task was to put our model on the robot. In order to have an integrated model, we needed the robot to perceive the world (via the CMVision system), give that information to the ACT-R model, allow the model to reason about the game and decide on a hiding place, have the robot go to the desired location, and then receive feedback (verbally).

At the beginning of a game, ACT-R sends a request to look for objects. The robot turns to look at the entire room, building a list of all of the recognized objects. Duplicate observations are removed based on object type and location, and the list is returned to the model. ACT-R uses its cognitive model with the object list to determine where it wants to hide. This hiding place is sent to the Wax system using the object record from the list and a relative location (e.g., under). The Wax system then uses simple geometry to apply the relative location to the object's observed position, from the robot's current viewpoint, using the object's a priori physical size. The resulting Cartesian coordinates of the hiding location are then sent to Trulla and the robot navigates to specific coordinates.

ACT-R is informed upon arrival at the hiding place and then asks the user for feedback on how well it hid. The user replies with natural speech one of a set of utterances that provide feedback on the quality of the hiding place and optionally, suggestions for the next time, such as "that object is too small to hide under". The speech is processed by the Nautilus speech understanding system, and the resulting encoded meaning sent to ACT-R. The cognitive model is updated to improve its decision-making and ACT-R tells the robot to go back to its starting position in preparation for the next game.

Parts of the interaction and robot behavior had to be changed from E's – for example, it is impossible for a robot to cover its eyes since it has no hands. When the robot wants to hide its “eyes” we simply have the robot turn 180 degrees from “It” (see Figure 2). A full set of movies of the robot hiding can be found at <http://www.aic.nrl.navy.mil/~trafton/hideseek.html>.



Figure 2: The robot turning 180 degrees to “close” its eyes.

Model and Robot behavior: Seeking

So far, we have shown a computational cognitive model that allows a mobile robot to hide in the same manner that a 3 1/2 year old child does. Our current system shows strong support for our object-relationship hypothesis about how children learn to play hide and seek, but we have not yet shown strong evidence for our representational hypothesis: that building a system that uses representations and processes similar to a person's will exhibit more natural behaviors. If this hypothesis is correct, we would expect to be able to use our existing system hiding system to seek for a person. The seeking system should exhibit several interesting behaviors. First, it should seek according to its own model of hiding. That is, it should search in places that it thinks are plausible for it to hide in.² Second, it should be able to deal with novel objects or objects that were not in its original environment. Third, it should be able to accomplish this seeking behavior without new learning mechanisms while using its current representations and algorithms. This seeking behavior would be strong evidence for our representational hypothesis: a system learning to hide and then using that information to search in places that would be "natural" for it to hide in.

In order to explore how our existing system would seek for a person after it had learned how to hide, we went through several straightforward steps. First, we ran the model as above, allowing it to learn different pertinent features of objects and object-relations. We then "froze" the model. In order to allow it to seek, we gave it two more pieces of information: (1) what a person "looked like" (e.g., the person would wear a blue shirt which was identifiable by

² Clearly, our robot can not bend or change shape like a young child. As a simplification for both the model and the robot, we assume that the hider is small (approximately child size) and does not contort itself a great deal or squeeze itself into a location that is rather smaller than itself is.

CMVision) and (2) how to start the game (e.g., a location to start from; what to count to, etc.).

In order to seek for a person, the computational cognitive model determined where it would hide and then gave those coordinates for the robot to look there. If it did not find the person in that location, it searched in the next place that it would hide until either it had found the person or it had run out of locations to search. We did not clear the model's "individual preferences" (e.g., locations that had higher or lower levels of activation); the model would search those locations in approximate (because of noise) order of activation. We also changed the environment slightly (i.e., added additional objects it already knew about, moved the location of other objects, etc.).

Both the model and robot behaved as expected. The robot systematically searched different locations that it had learned were acceptable hiding places until it found the person hiding. Over multiple games, it searched locations in different orders. Importantly, it did not attempt to search for a person in locations that would have been very "odd." For example, while it could have found a person hiding out in the open, it did not systematically search all the open space for a person hiding out in the open. Instead, it searched where it thought it would have hidden. A full set of movies of the robot seeking can be found at <http://www.aic.nrl.navy.mil/~trafton/hideseek.html>.

Conclusions

This paper suggested two different hypotheses, a specific object-relationship hypothesis dealing with how children learn to play hide and seek, and the second representational hypothesis dealing with the types of representations and algorithms or procedures that should be used for intelligent systems. Both hypotheses were supported. The object-relationship hypothesis was that children learn how to play a credible game of hide and seek not by using spatial perspective

taking but by learning the features and relations of objects (e.g., hiding inside of something is usually a good hiding place). This hypothesis was supported by both empirical and computational evidence. The case study of E showed that a 3 1/2 year old child who did not have perspective-taking skills was able to learn how to play a credible game of hide and seek. She did this primarily by learning to hide inside and under different objects. Importantly, she exhibited almost no spatial perspective taking in her hiding behavior. We also supported our object-relationship hypothesis by building a computational cognitive model in ACT-R. Our cognitive model matches E's hiding behavior at a qualitative level and makes the same type of errors that E made. Clearly, there is a limit to what can be learned using this type of hiding behavior – hiding behind things can not be done, locations are used multiple times, etc. The learning mechanisms we used in our model are quite general and used in other cognitive models, so we did not need to invent any new learning methods. Finally, we put our model on a physical robot to embody the computational model.

We also proposed and supported a representational-level hypothesis. Our hypothesis was that building a computational or robotic system that uses representations and processes similar to a person's will be able to work very well with a person because less conversion will be needed to translate between different representations. We supported this hypothesis by taking the “hiding” model and applied it to seeking. The model successfully searched for a person using the same representations and processes that it had learned and used while learning how to hide. Clearly, our approach could lead the system to make systematic errors: it would not expect a person to hide on the ceiling or search for very small people, etc. It would not use perspective taking for seeking or even assume that the hider would move locations. However, it could be

argued that this is exactly the type of seeking behavior that a child would exhibit. Further empirical studies could be performed to examine this exact issue.

Integrating a computational cognitive model with a robotic system gave us several advantages. First, our system allowed the cognitive model to do the “thinking and reasoning” aspects of the task and the robot’s low-level mobility code to do the navigation and wayfinding. This separation between high (cognitive model) and low (mobility) code seems like a natural dividing point for what computational cognitive models are good at (thinking, reasoning, problem solving, etc.) and what more engineering models are good at (low level perceptual issues, navigation, search, etc.). Finally, by putting our cognitive model on our robot, we have made a large step to embodied cognition.

Acknowledgements

This work was partially supported by the Office of Naval Research to the first two authors under work order number N0001402WX20374. Additional funding was supplied by DARPA IPTO under the MARS program. The views and conclusions contained in this document should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U. S. Navy.

References

- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, 26(1), 39-83.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (2002). Spanning seven orders of magnitude: A challenge for cognitive modeling. *Cognitive Science*, 26.
- Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, 38, 341-380.
- Anderson, J. R., Conrad, F. G., & Corbett, A. T. (1989). Skill acquisition and the LISP tutor. *Cognitive Science*, 13, 467-505.
- Anderson, J. R., & Lebiere, C. (1998). *Atomic components of thought*. Mahwah, NJ: Erlbaum.
- Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, 96(4), 703-719.
- Anzai, Y., Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86(2), 124-140.
- Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: Further evidence. *Child Development*, 62, 1227-1246.
- Baillargeon, R., Spelke, E., & Wasserman, S. (1985). Object permanence in 5 month old infants. *Cognition*, 20, 191-208.
- Borenstein, J., & Koren, Y. (1991). The Vector Field Histogram - fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, 7(3), 278-288.
- Bower, T. G., Broughton, J., & Moore, M. K. (1971). Development of the object concept as manifested in changes in the tracking behavior of infants between 7 and 20 weeks of age. *Journal of Experimental Child Psychology*, 11(2), 182-193.

- Bruce, J., Balch, T., & Veloso, M. (2000). Fast and inexpensive color image segmentation for interactive robots. In *Proc. of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Vol. 3, pp. 2061-2066). Takamatsu, Japan.
- Bugajska, M. D., Schultz, A. C., Trafton, J. G., Mintz, F. E., & Gittens, S. (2001). Building adaptive computer generated forces: The effect of increasing task reactivity on human and machine control abilities. In *Late-Breaking Papers at the 2001 Genetic and Evolutionary Computation Conference (GECCO 2001)*. San Francisco, CA.
- Bugajska, M. D., Schultz, A. C., Trafton, J. G., Taylor, M., & Mintz, F. E. (2002). A hybrid cognitive-reactive multi-agent controller. In *Proceedings of 2002 IEEE/RSJ International conference on Intelligent Robots and Systems (IROS 2002)*. Switzerland.
- Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In J. R. Anderson & C. Lebiere (Eds.), *Atomic Components of thought* (pp. 167-200). Mahwah, NJ: Lawrence Erlbaum.
- Carbonell, J. G., Knoblock, C. A., & Minton, S. (1991). Prodigy: An integrated architecture for planning and learning. In K. V. Lehn (Ed.), *Architectures for intelligence* (pp. 241-278). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Flavell, J. H., Omonson, R. C., & Latham, C. (1978). Solving spatial perspective taking problems by rule vs. computation: A developmental study. *Developmental Psychology*, 14, 462-473.
- Harrison, A. M., & Schunn, C. D. (2002). ACT-R/S: A computational and neurologically inspired model of spatial reasoning. In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the Twenty Fourth Annual Meeting of the Cognitive Science Society* (pp. 1008). Fairfax, VA: Lawrence Erlbaum Associates.
- Harrison, A. M., & Schunn, C. D. (2003a). ACT-R/S: Look Ma, No "Cognitive-map"! In *International Conference on Cognitive Modeling*.
- Harrison, A. M., & Schunn, C. D. (2003b). Segmented spaces: Coordinated perception of space in ACT-R. In F. Detje, D. Dorner & H. Schaub (Eds.), *The logic of cognitive systems: Proceedings of the fifth international conference on cognitive modeling* (pp. 307). Bamberg, Germany.

- Hughes, K., Tokuta, A., & Ranganathan, N. (1992). Trulla: An algorithm for path planning among weighted regions by localized propagations. In *Proc. of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 469-476). Raleigh, NC: IEEE Press.
- Huttenlocher, J., & Kubicek, L. (1979). The coding and transformation of spatial information. *Cognitive Psychology*, 11, 375-394.
- Kleeman, J. (1973). The Peek-A-Boo game: Its evolution and associated behavior, especially bye-bye and shame expression during the second year. *Journal of the American Academy of Child Psychiatry*, 12(1), 1-23.
- Lee, F. J., & Gamard, S. J. (2003). Hide and seek: Using computational cognitive models to develop and test autonomous cognitive agents for complex dynamic tasks. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. Boston, MA.
- Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory: Cross-task prediction of performance in ACT-R. *Cognitive Systems Research*, 1, 99-118.
- Moravec, H. P., & Elfes, A. E. (1985). High resolution maps from wide angle sonar. In *Proc. of the IEEE International Conference on Robotics and Automation* (pp. 116-121). St. Louis, MO: IEEE Press.
- Newcombe, N., & Huttenlocher, J. (1992). Children's early ability to solve perspective taking problems. *Developmental Psychology*, 28, 654-664.
- Perzanowski, D., Schultz, A., & Adams, W. (1998). Integrating Natural Language and Gesture in a Robotics Domain. In *Proceedings of the IEEE International Symposium on Intelligent Control: ISIC/CIRA/ISAS Joint Conference* (pp. 247-252). Gaithersburg, MD: National Institute of Standards and Technology.
- Perzanowski, D., Schultz, A., Adams, W., Bugajska, M., Marsh, E., Trafton, J. G., et al. (2002). Communicating with teams of cooperative robots. In A. C. Schultz & L. E. Parker (Eds.), *Multi-Robot Systems: From Swarms to Intelligent Automata* (pp. 16-20). The Netherlands: Kluwer.

- Perzanowski, D., Schultz, A., Adams, W., & Marsh, E. (2000). Using a Natural Language and Gesture Interface for Unmanned Vehicles. In G. R. Gerhart, R. W. Gunderson & C. M. Shoemaker (Eds.), *Proceedings of the Society of Photo-Optical Instrumentation Engineers* (Vol. 4024, pp. 341-347).
- Perzanowski, D., Schultz, A., Adams, W., Marsh, E., & Bugajska, M. (2001). Building a multimodal Human-Robot Interface. *IEEE Intelligent Systems*, 16-20.
- Petty, M. D. (2001). Do we really want computer generated forces that learn? In *Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation (CGF&BR)*. Norfolk, VA.
- Piaget, J. (1954). *The construction of reality in the child*. New York, NY: Basic Books.
- Piaget, J., & Inhelder, B. (1948). *The child's conception of space*. London: Routledge and Kegan Paul.
- Previc, F. H. (1998). The neuropsychology of 3-D space. *Psychological Bulletin*, 124(2), 123-164.
- Schultz, A., & Adams, W. (1998). Continuous localization using evidence grids. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation* (pp. 2833-2839). Leuven, Belgium: IEEE Press.
- Schultz, A., Adams, W., & Yamauchi, B. (1999). Integrating exploration, localization, navigation and planning with a common representation. *Autonomous Robots*, 6, 293-308.
- Shepard, R., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701-703.
- Skubic, M., Perzanowski, D., Blisard, S., Schultz, A., & Adams, W. (in press). Spatial Language for Human-Robot Dialogs. *IEEE Transactions on Systems, Man, and Cybernetics*.
- Smith, L. B. (2000). From knowledge to knowing: Real progress in the study of infant categorization. *Infancy*, 1, 91-98.
- Taylor, H. A., & Tversky, B. (1992). Spatial mental models derived from survey and route descriptions. *Journal of Memory & Language*, 31, 261-292.

- Trafton, J. G., Kirschenbaum, S. S., Tsui, T. L., Miyamoto, R. T., Ballas, J. A., & Raymond, P. D. (2000). Turning pictures into numbers: Extracting and generating information from complex visualizations. *International Journal of Human Computer Studies*, 53(5), 827-850.
- Trafton, J. G., Schultz, A. C., Bugajska, M. D., Gittens, S., & Mintz, F. E. (2001). An investigation of how humans and machines deal with increases in reactivity. In *The proceedings of the tenth conference on Computer Generated Forces and Behavioral Representation (CGF&BR)*. Norfolk, VA.
- Trickett, S. B., Ratwani, R. M., & Trafton, J. G. (under review). Real-World Graph Comprehension: High-Level Questions, Complex Graphs, and Spatial Cognition. *Psychological Bulletin*.
- VanLehn, K. (1991). Rule acquisition events in the discovery of problem-solving strategies. *Cognitive Science*, 15, 1-47.
- Wallace, R., Allan, K. L., & Tribol, C. T. (2001). Spatial perspective-taking errors in children. *Perceptual and Motor Skills*, 92(3), 633-639.
- Yamauchi, B., Schultz, A., & Adams, W. (1998). Mobile Robot Exploration and Map Building with Continuous Localization. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation* (pp. 3715-3720). Leuven, Belgium: IEEE.